

Our Ref.: 2380-207

U.S. PATENT APPLICATION

Inventors: Vesa LEHTOVIRTA
Alain MAUPIN

Invention: METHOD AND APPARATUS FOR RELEASING CONNECTIONS IN
AN ACCESS NETWORK

NIXON & VANDERHYE P.C.
ATTORNEYS AT LAW
1100 NORTH GLEBE ROAD
8TH FLOOR
ARLINGTON, VIRGINIA 22201-4714
(703) 816-4000
Facsimile (703) 816-4100

SPECIFICATION

METHOD AND APPARATUS FOR RELEASING CONNECTIONS IN AN ACCESS NETWORK

RELATED PRIORITY APPLICATIONS

This application claims priority from commonly-assigned U.S. Provisional
5 Patent Application Serial No. 60/184,327, filed on February 23, 2000 and commonly-assigned
U.S. Provisional Patent Application Serial No. 60/183,325, filed on February 18, 2000, both of
which are entitled "Method and Apparatus for Releasing Radio Access Bearers In A Radio
Access Network," the disclosures of which are incorporated herein by reference.

FIELD OF THE INVENTION

10 The present invention relates to radio communication networks and other
communication networks, where a subscriber unit, e.g. a mobile radio station, may employ
several simultaneous connections to the network.

BACKGROUND AND SUMMARY OF THE INVENTION

15 Mobile communications have developed from first generation, analog-based
mobile radio systems to second generation digital systems, such as the European Global
System for Mobile communications (GSM). Current developments for a third generation of
mobile radio communication include the Universal Mobile Telephone communications System
(UMTS) and the IMT 2000 system. For simplicity, third generation systems are referred to
simply as UMTS. In simple terms, UMTS is "communication to everyone, everywhere,"
20 where communication also includes providing information using different types of media, i.e.,
multimedia communications.

25 Second generation mobile/cellular telecommunications networks are typically
designed to connect and function with Public Switched Telephone Networks (PSTNs) and
Integrated Services Digital Networks (ISDNs). Both of these networks are circuit-switched
networks (rather than packet-switched) and handle relatively narrow bandwidth traffic.
However, packet-switched networks, such as the Internet, are very much in demand and
handle much wider bandwidth traffic than circuit-switched networks. While wireline
communication terminals, e.g., personal computers, are capable of utilizing the wider packet-
switched network bandwidth, wireless user equipment units (UEs) are at a considerable

disadvantage because of the limited bandwidth of the radio/air interface between UEs and packet-switched networks.

UEs are currently limited in the data rates for data communications services such as facsimile, electronic mail, and Internet. The demand is growing for higher data transfer speeds in order the “surf the net” using UEs with fast access to text, images, and sound. Multimedia applications demand high peak bit rates in short bursts, particularly when information is downloaded to the UE. Another challenging multimedia UE application is simultaneous voice and data, e.g., PC application sharing or shared whiteboard. Although this latter type of multimedia application may not require particularly high bit rates, it does require real time, continuous operation because of the voice content. A demanding circuit-switched application (rather than packet-switched as in the Internet application) requiring relatively high bit rates is video conferencing. In order for mobile video conferencing to become practical, the amount of user bandwidth required must be reduced to a minimum without sacrificing image quality.

A UMTS Wideband-Code Division Multiple Access (WCDMA) radio access network provides wireless access at very high data rates and supports enhanced services not realistically attainable with the first and second generation mobile communication systems. WCDMA currently supports 5 MHz-15 MHz, and in the future, promises an even greater bandwidth. In addition to wide bandwidth, WCDMA also improves the quality of service by providing robust operation in fading environments and transparent (“soft”) handoffs between base stations. Multiplath fading is used to advantage to enhance quality, i.e., using a RAKE receiver and improved signal processing techniques, contrasted in narrowband systems where fading substantially degrades signal quality.

A UMTS Terrestrial Radio Access Network (UTRAN) responds to radio access service requests by allocating resources needed to support a communication with a UE. The UTRAN includes plural base stations for communicating with UEs over a radio air interface using radio channel resources allocated by a radio network controller connected to the base stations. External network service nodes that interface with external networks, communicate with UEs via the UTRAN. When one of the service nodes requires communication with a UE, the service node requests a radio access “bearer” (RAB) from the UTRAN rather than a

specific radio channel resource. A radio access bearer is a logical connection with the UE through the UTRAN and over the radio air interface and corresponds to a single data stream. For example, one radio access bearer may support a speech connection, another bearer may support a video connection, and a third bearer may support a data packet connection. Each radio access bearer is associated with quality of service (QoS) parameters describing how the UTRAN should handle the data stream. Examples of quality of service parameters include data rate, variability of data rate, amount and variability of delay, guaranteed vs. best effort delivery, error rate, etc. Although the term “radio access bearer” is sometimes used for purposes of the following description, the invention applies to any type of “connection,” and is not limited to logical connections like RABs, a particular type of physical connection, etc.

Radio access bearers are dynamically assigned to UTRAN transport and radio channel resources by the UTRAN. The radio access bearer service and the UTRAN isolate the details of transport and radio resource allocation handling as well as details of radio control, e.g., soft handoff. The UTRAN approach is different from traditional approaches where an external network and/or an external network service node is involved in the details of requesting, allocating, and controlling specific radio connections to and from the mobile radio. Instead, the external network service node only needs to request a radio access bearer service over a RAN interface to the UTRAN along with a specific quality of service for a communication to a specific mobile radio. The UTRAN provides the requested service at the requested quality of service (if possible).

Plural radio access bearers may be established and released independently to one UE including bearers from different networks. Moreover, plural radio access bearers, e.g., one carrying circuit-switched information and another carrying packet-switched information, intended for the specific UE may be multiplexed onto the same CDMA channel. Each bearer may have its own transport connection through the UTRAN, or it may be multiplexed with other bearers onto one transport connection.

To initiate a radio access bearer service, a request is transmitted to the UTRAN for communication with a UE. One or more parameters accompany the radio access bearer service request. When establishing each bearer, the UTRAN “maps” or allocates the radio access bearer to physical transport and radio channel resources through the UTRAN and over

the radio air interface, respectively. The transport connection between nodes in the UTRAN may be for example an ATM type connection. A radio channel over the air interface includes one or more CDMA spreading codes.

The mapping is based on the one or more parameters associated with the radio access bearer service request. In addition to quality of service parameters, the parameters may also include one or more traffic condition parameters like a congestion level on a common channel, an interference level in the geographic location area in which the UE is currently operating, a distance between the UE and the base station, radio transmit power, the availability of dedicated channel resources, the existence of a dedicated channel to a UE, and other traffic parameters or conditions.

An example Universal Mobile Telecommunications System (UMTS) 10 is shown in Fig. 1. A representative, core network 16, includes a circuit-switched core network (CS CN), shown as box 18, and a packet-switched core network (PS CN), shown as box 20. The circuit-switched core network includes nodes, e.g., Mobile Switching Centers (MSC) 18, Home Locations Register (HLR), Gateway MISC (GMSC), etc., that provide circuit-switched services. The packet-switched core network includes nodes, e.g., Serving GPRS Support Nodes (SGSN) 20, Gateway GPRS Support Node (GGSN), HLR, etc., that are tailored to provide packet-switched services. The CSCN 18 is coupled to an external circuit-switched network 12, such as the Public Switched Telephone Network (PSTN) or the Integrated Services Digital Network (ISDN). The packet-switched core network 20 is coupled to an external packet-switched network 14, such as the Internet.

Each of the core networks 18 and 20 is coupled to a UMTS Terrestrial Radio Access Network (UTRAN) 22 that includes one or more Radio Network Controllers (RNCs) 26. Each RNC is coupled to a plurality of base stations (BSs) 28 and to other RNCs in the UTRAN 22. Each base station 28 corresponds to one access point (one sector or cell) or includes plural access points. Radio communications between one or more base station access points and wireless user equipment unit (UE) 30 are by way of a radio interface. Radio access in this non-limiting example is based on Wideband-CDMA (W-CDMA) with individual radio channels distinguished using spreading codes. Wideband-CDMA provides wide radio

bandwidth for multi-media services including packet data applications that have high data rate/bandwidth requirements.

Fig. 2 illustrates an example where a UE has four simultaneous radio access bearers (RABs) with core networks via the UTRAN: one RAB towards the CSCN 18, and three RABs towards the PSCN 20. For simplicity only a single RNC 26 is shown in Figure 2. In this example, RAB#1 could be used for a speech call, RAB#2 for web browsing, RAB#3 for downloading files with file transfer protocol (FTP), and RAB#4 for sending electronic mail.

UMTS network nodes, such as BSs, RNCs, MSCs, GPRS nodes, etc., may employ a modular and distributed architecture where several processor boards are coupled to a switch. Referring to the generic node 40 in Fig. 3, there may be several processors 44 on each processor board 42. The processors 44 communicate via the switch 46. The switch could be, for example, an ATM-type switch. Further details of such an architecture are described, for example, in commonly-assigned, co-pending application Serial No. 09/039,453 entitled, "Asynchronous Transfer Mode Platform for Mobile Communications," filed on March 16, 1998, the disclosure of which is incorporated herein by reference. Each processor 44 may include a self-detecting failure mechanism with one or more hardware detectors and/or software error detection algorithms.

When a connection, such as a radio access bearer (RAB), is established through a network node, a processor is allocated to handle the connection. Fig. 4 shows how four RABs of a UE 30 could be allocated to different processors in an RNC node 26 and a PSCN node 20. For simplicity, only the RABs of one UE 30 are shown, and the base station node is omitted. RAB#1 is handled by processor 1 on processor board C and by CS CN (for which details are omitted). RABs#2 is handled by processor 2 on processor board C and by processor 1 on processor board A. RAB#3 is handled by processor 2 on processor board C and by processor 3 on processor board B. RAB#4 is handled by processor 4 on processor board D and by processor 4 on processor board B.

It is not uncommon for failures of some sort to occur in a node. The whole node or only a part of the node may fail. If the whole node fails, all connections through the

node are lost. Consider the example depicted in Fig. 4 where there is a complete failure of the PSCN node 20. As a result, RAB#2, RAB#3, and RAB#4 are lost. However, RAB#1 survives because it is not using the failed PSCN node 20. A partial failure may affect only one device or board, while other devices or boards remain fully operational. An example of a partial failure is when a processor 44 or a processor board 42 crashes or is restarted. In Fig. 4, if processor board C in the RNC 26 fails, RAB#1, RAB#2, and RAB#3 are lost. However, RAB#4 survives because RAB#4 is not supported by the failed processor board C. If processor 3 in board B in the PSCN node 20 experiences a restart, RAB#3 is lost, but other RABs supported by board B survive, including RAB#4 to processor 4.

When a connection is lost in a node due to a failure of some sort, other nodes assigned to support that connection may not detect that the connection has actually been lost. Unless those other nodes are informed, the unreleased connection and associated supporting resources remain reserved for the connection, even though they are not being used. In the example shown in Fig. 5, if processor board C in the RNC 26 fails, RAB#1, RAB#2, and RAB#3 are lost. However, RAB#4 survives because RAB#4 is not using the failed processor board C. The RNC 26 should therefore instruct the CSCN 18 to release RAB#1. The RNC should also instruct the PSCN 20 to release RAB #2 and RAB #3. However, RAB#4 need not be released because it was not affected by the failure. In fact, RAB#4 should not be released if it carries a service independent of the other RABs.

Accordingly, it is desirable to selectively release resources affected by a partial failure in a node, while allowing unaffected node element, connections, and resources to remain intact and functioning. There are different approaches to achieving these ends. Preferably, those approaches should be easily implemented in existing systems, e.g., using messages already-defined by or consistent with UTRAN standard signaling protocols.

The present invention meets the above-identified objectives. Initially, communication connections are established between an external network and subscriber units (e.g., wireline telephones, wireless UEs, etc.) by way of an access network. The networks may provide wireline service, wireless service, or both. As described earlier, a connection includes any type of logical or physical communications connection that corresponds to a single information stream. A subscriber unit may employ one or plural communication connections.

The subscriber unit connection is supported by plural nodes. When a failure is detected in a node, those subscriber unit connections affected by that failure are determined. A failure may include a complete failure in a node, a partial failure in a node, failure in one of several devices in a node, or any other incident that would impact the ability of that node to support a subscriber connection. A message identifying those affected subscriber unit connections is sent to one or more other nodes. Affected subscriber unit connections identified in the message are released. However, those unaffected connections not included in the message are maintained.

In an example embodiment, a list is generated that identifies the subscriber units affected by the detected failure along with the connections affected by the failure. This list is included in the message sent to one or more other nodes supporting a connection of identified subscriber unit connections. In the example context of a radio communications system like UMTS, the node in which the failure is detected may be any one of an external network node, a core network node, a radio access network node such as an RNC or a base station, or a UE.

In another example implementation of the present invention, network addresses, e.g., IP addresses, are assigned to devices in the node. When a subscriber unit connection is established, an address for each device associated with the subscriber unit connection is sent to other nodes. If a failure is detected in one of the devices, a message including the network address of the failed device is sent to one or more other nodes. For nodes containing plural processor boards coupled by a switch, where each processor board includes plural processors, such a message may identify the addresses of the plural processor on a failed board. As a result, the node(s) receiving the message release subscriber unit connections associated with that failed processor board.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the following description of preferred, non-limiting example embodiments, as well as illustrated in the accompanying drawings. The drawings are not necessarily to scale

or comprehensive, emphasis instead being placed upon illustrating the principles of the invention.

Fig. 1 is a block diagram of a UMTS-type system in which the present invention may be advantageously employed;

5 Fig. 2 is a diagram illustrating four connections from two core network nodes to a UE based on the UMTS system shown in Fig. 1;

Fig. 3 shows a modular and distributed node architecture that may be employed in one or more of the nodes shown in Fig. 1;

10 Fig. 4 is an illustration of how four connections to a UE are supported using the node architecture shown in Fig. 3;

Fig. 5 illustrates an example where one of the processor boards in a node fails, and how that failure impacts connections to the UE;

Figs. 6, 7 and 8 are examples of a failed processor board in which a message is sent from the affected node as a result of that failure;

15 Fig. 9 is a flowchart diagram illustrating procedures in accordance with an example implementation of the invention when an individual processor fails;

Fig. 10 is a flowchart diagram illustrating procedures in accordance with an example implementation where a node fails;

20 Fig. 11 is a function block diagram illustrating an example where a processor board fails;

Fig. 12 is an example implementation of the present invention that employs IP addresses of node elements;

Fig. 13 is a flowchart diagram illustrating example procedures for implementing the example implementation shown in Fig. 12; and

Fig. 14 is an example implementation of the present invention for a failure effected in a UE.

DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, for purposes of explanation and not limitation, specific details are set forth, such as particular embodiments, procedures, techniques, etc., in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that the present invention may be practiced in other embodiments that depart from these specific details. For example, the following description is in the context of a downlink example from the radio network to wireless user equipment. Those skilled in the art will appreciate that the present invention may also be implemented in the opposite, uplink direction. Moreover, while the present invention finds advantageous application to and is described in the context of a radio communications environment, the invention may also be used in a wireline network environment. In such a wireline network application, the subscriber units would be wireline subscriber units and the access network would be a wireline access network interfacing the subscriber units and other networks such as Internet Service Provider (ISP) networks, the Internet, the PSTN, the ISDN, etc.

In some instances, detailed descriptions of well-known methods, interfaces, devices and signaling techniques are omitted so as not to obscure the description of the present invention with unnecessary detail. Moreover, individual function blocks are shown in some of the figures. Those skilled in the art will appreciate that the functions may be implemented using individual hardware circuits, using software functioning in conjunction with a suitably programmed digital microprocessor or general purpose computer, using an Application Specific Integrated Circuit (ASIC), and/or using one or more Digital Signal Processors (DSPs).

A first approach to a node failure is to send a RESET message which indicates that all connections associated with that failed node should be released. An example of such a RESET message is a UTRAN Iu Interface RANAP RESET message as defined by the 3GTS 25.413 UTRAN standard. Fig. 6 shows an example of this general reset approach. All RABs associated with the failed PSCN node 20 are released, even though only one RAB is actually

affected by the failure. More specifically, after processor board A fails in the PSCN node, a RESET message is sent to the RNC 26. The RESET message indicates that the PSCN node 20 has had a node-wide failure, and as a result, the RNC 26 should release resources (including all RABs and signaling connections) associated with the PSCN node 20. However, the failure in board A only actually impacts one RAB to UE1. Unfortunately, five other RABs to UE1 and UE2, which are not affected by the processor board A failure, are also released. Only the RAB between UE1 and the CSCN remains intact. Since there may be many boards in a node, many RABs would be released for no reason in this kind of node-wide failure message.

A second approach is to release each affected RAB individually using control signaling, e.g., existing RANAP control signaling. In the example shown in Fig. 7, a processor board B fails, and the RABs marked with dotted lines are lost. Individual RAB Release request messages are sent by the PSCN 20 to the RNC 26 to each UE having a RAB that uses the failed processor board C. A “Release RAB5” message is sent to UE2, and a “Release RAB6 and RAB7” message is sent to UE3. The benefit of this approach is that unaffected RABs are maintained. A drawback is that since there may be many RABs affected by the failed processor board C, a large and immediate signaling burden is placed on the radio and core networks.

The third approach is to just release affected UEs. Consider the situation illustrated in Fig. 8. When a partial failure, e.g., a processor restart, is detected in processor 3 of processor board B in the PSCN node 20, a single message is sent to the RNC 26 listing all UEs having at least one RAB supported by the failed processor 3. An example of such a message is RESET RESOURCE message in the RANAP protocol. Fig. 8 shows an example: RESET RESOURCE UE2, UE3. When this message is received in the RNC 26, the RNC 26 should release all resources (RABs and signaling connections) of all listed UEs. While this reduces the signaling load, a disadvantage of this affected-UE approach is that unaffected RABs for those UEs are released unnecessarily.

None of these three approaches maintain RABs unaffected by a partial failure without causing a significant “release” signaling load. The “Failure in Node” routine outlined in the flowchart of Fig. 9 overcomes these shortcomings. A failure is detected in a network

node (step 100). A network failure includes any type of failure that affects the node's ability to support or otherwise render service to a particular connection to a user equipment. Of course, such failures include situations where an entire node is disabled, but they also include any partial failure in the node.

5 Assuming the node is configured for example as shown in Fig. 3, one of the processors may detect an error, and therefore, must perform some type of processor restart or error handling operation. As mentioned above, such processors and other devices often have self-detecting mechanisms that detect failures or errors, and send appropriate messages based on the type of failure or error detected. Some other node controller (not shown) may also be employed.

10 A decision is made in the node whether the failure is a partial failure (block 102). If not, it may be appropriate in some circumstances to assume that the entire node has failed and must be reset. Accordingly, the node sends a reset or similar type message to one or more other nodes such as RNCs, base stations, UEs, and other core network nodes (block 104). If a partial failure is detected in block 102, e.g., one of plural processors is restarted, the node sends a message to one or more nodes that selectively indicates which connections are affected by the failure and should be released. Those connections unaffected by the partial failure are maintained (block 106). Using this approach, unaffected connections are not released unnecessarily. On the other hand, there is no need to send individual
15 messages for each connection that is affected by the partial failure, which keeps the control signaling load to a manageable level.

Fig. 10 employs a specific example of such a message. As with Fig. 9, a failure is detected in a network node (block 110), and a decision is made whether the failure is a partial or complete failure (block 112). If the failure is node-wide, a node reset message is sent
20 to one or more nodes (block 114). On the other hand, for a partial failure, the node forms a list of UEs and connections (in this case radio access bearers) specifically affected by the partial failure. The list could include for example UE identifiers, (e.g., UE's IMSI), and RAB identifiers (block 116). The list is then sent to one or more nodes (block 118).

The list may be sent in an existing message or a new message may be introduced to indicate a partial failure in the node. The one or more nodes that receive the message, release all radio access bearers (RABs) on the list, i.e., all RABs associated with the UE connection, but maintain those RABs to that node that are not listed (block 120). In one variation, if all radio access bearers of a user equipment are affected, no RAB IDs need be sent. Instead, the message sent to the node(s) simply indicates that all associated RABs for a particular UE should be released. In this latter instance, there may be situations where it is desirable to maintain a signaling connection to that user equipment unit, e.g., registration, unless the message specifically indicates that the signaling connection should be released as well.

Consider the example shown in Fig. 11 where a processor board B fails. Since RABs in UE2 (RAB5) and UE3 (i.e., RAB6 and RAB7) are affected by that failure, the PSCN 20 sends a message, e.g., an "N_RESET" message, to the RNC listing these RABs and the IDs of the affected UEs. The RNC 26 then releases these RABs. Since both RABs of UE3, (i.e. RAB6 and RAB7), are affected, the RAB IDs need not be sent to the RNC to be released. The N_RESET message simply indicates that all RABs for UE3 are to be released. Because RAB3 and RAB4 of UE2 are not affected, they are not released. RAB1 and RAB2 for UE1 also remain intact. The N_RESET message is shown only being sent to the RNC for simplicity; however, that message may be sent to one or more base stations, UEs, and/or other core network nodes.

Another example and non-limiting implementation of the invention uses an IP address of the failed device. In IP-capable nodes, a separate IP address may be assigned for each device, e.g. a processor, that handles packet data traffic. When an RAB is set up, the IP address of the allocated device is sent to one or more peer node(s). The peer node(s) then routes the traffic of that to this IP address.

Referring to the example shown in Fig. 12, if the processor (IP address 39) fails, the PSCN node 20 sends a message to the RNC 26 with the IP address of the failed processor. The RNC then releases all RABs associated with this IP address. Again, this single message, selective RAB release approach is efficient because of its low signaling load and effective because RABs are not unnecessarily released.

There is typically a signaling connection for each UE between the RNC and UE, the RNC and the PSCN, and the RNC and the CSCN, (e.g., SCCP connections). In accordance with another aspect of the present invention, the connection release message may also include an indication whether the UE signaling connection should be released when an error or a failure is detected that affects the UE. This aspect of the invention may be useful in situations when a fault or error affecting the UE requires all of its RABS to be released, but where it is desirable to maintain the signaling connection between the UE and the radio network. One reason why the signaling connection may be maintained is to permit the UE to update the radio network with its current location. Accordingly, for each UE in the list of affected UEs in the message, an associated signaling connection indicator such as a flag may be set (release the signaling connection) or reset (maintain the signaling connection) in addition to the list of affected RABS. If a list of RABS to be released is included in the message, and there are unlisted RABS to be maintained, the signaling flag is ignored because the UE signaling connection must be maintained for unaffected RABS.

A "Failure in IP-Node" routine is now described in conjunction with the flowchart illustrated in Fig. 13. An IP address is assigned to each device in a node that handles data packets (block 130). When a radio access bearer is set up through that node, the node sends the IP addresses of those devices that have been allocated to support the radio access bearer to one or more nodes (block 132). A decision is made whether one of those allocated devices has failed (block 134). If it has, the IP-node sends a message to one or more nodes with the address of the failed device(s) (block 138). The one or more nodes, upon receiving that message, releases all radio access bearers associated with this IP address. However, as just explained, a signaling connection may be maintained if desired or otherwise appropriate (block 140). Whether to maintain the signaling connection for a UE may be indicated with a list of those UEs whose signaling connection is to be released. The signaling connection is maintained for the UEs that are not included in the list.

If a device failure is not detected, a decision is made whether a board in the node, e.g., a processor board, has failed (block 136). If it has, the IP node sends a message to one or more nodes listing the IP addresses of all IP devices that are located on that failed board (block 142). Upon receiving that message, the one or more nodes release all radio

access bearers associated with the IP addresses listed in the message. Again, the signaling connection may be maintained (block 144). Whether to maintain the signaling connection for a UE may be indicated with a list of those UEs whose signaling connection is to be released. The signaling connection is maintained for the UEs that are not included in the list.

5 The present invention may be applied to failures detected in a wireless user equipment unit (UE). Fig. 14 illustrates in flowchart format a "Failure in UE" routine. A decision is made in block 150 whether a failure has occurred in the UE. If it has, the UE sends a message indicating whether the signaling connection is to be maintained with the UE (block 152). One or more nodes release all radio access bearers associated with the UE, but
10 maintain the signaling link with the UE, if so indicated (block 154).

 Thus, in the present invention, if a failure occurs in a node, unaffected connections are maintained, thereby improving the service provided to those users with unaffected connections. Moreover, a large amount of network release signaling is avoided. Of course, the invention is not restricted to the UMTS system and nodes shown in Fig. 1. Indeed, the invention can be employed in any radio communication system that employs
15 plural connections to a communications station. Other non-limiting applications include: the CDMA 2000 system, the GPRS system with SGSN and GGSN nodes, interactions between two routers, etc. In the UMTS system of Fig. 1, the invention may be implemented for interactions between RNC and RBS nodes, between two RNC nodes, between two RBS
20 nodes, between RBS and UE nodes, and between RNC and core network nodes, and between core network nodes.

 As explained earlier, the invention may be applied to wireline access network environments as well as radio access network environments. In a wireline communication system, connections are established between an external network and wireline subscriber units
25 (like PCs) by way of a wireline access network, such as a ISP. A failure is detected in a node in the wireline access network, and one or more subscriber unit connections affected by the detected failure are identified. A message identifying the one or more affected wireline subscriber unit connections (or just the affected wireline subscriber units) is sent so that one or more affected wireline subscriber unit connections may be released. Other aspects of the

invention described above, e.g., the list structures included in the message, may also be readily applied in this environment.

While the present invention has been described with respect to particular example embodiments, those skilled in the art will recognize that the present invention is not limited to those specific embodiments described and illustrated herein. Different formats, 5 embodiments, adaptations besides those shown and described, as well as many modifications, variations and equivalent arrangements may also be used to implement the invention. Therefore, while the present invention is described in relation to a preferred example embodiment, it is to be understood that this disclosure is only illustrative and exemplary of the present invention. Accordingly, it is intended that the invention be limited only by the claims 10 appended hereto.

09778960-000004